When silver glitters more than gold: Bootstrapping an Italian part-of-speech tagger for Twitter

Barbara Plank
University of Groningen
The Netherlands
b.plank@rug.nl

Malvina Nissim
University of Groningen
The Netherlands
m.nissim@rug.nl

arXiv:1611.03057v1 [cs.CL] 9 Nov 2016

Abstract

English. We bootstrap a state-of-the-art part-of-speech tagger to tag Italian Twitter data, in the context of the Evalita 2016 PoSTWITA shared task. We show that training the tagger on native Twitter data enriched with little amounts of specifically selected gold data and additional silver-labelled data scraped from Facebook, yields better results than using large amounts of manually annotated data from a mix of genres.

Italiano. *Nell'ambito della campagna di valutazione PoSTWITA di Evalita 2016, addestriamo due modelli che differiscono nel grado di supervisione in fase di training. Il modello addestrato con due cicli di bootstrapping usando post da Facebook, e che quindi impara anche da etichette "silver", ha una performance superiore alla versione supervisionata che usa solo dati annotati manualmente. Discutiamo l'importanza della scelta dei dati di training e development.*

1 Introduction

The emergence and abundance of social media texts has prompted the urge to develop tools that are able to process language which is often non-conventional, both in terms of lexicon as well as grammar. Indeed, models trained on standard newswire data heavily suffer when used on data from a different language variety, especially Twitter (McClosky et al., 2010; Foster et al., 2011; Gimpel et al., 2011; Plank, 2016).

As a way to equip microblog processing with efficient tools, two ways of developing Twitter-compliant models have been explored. One option is to transform Twitter language back to what pre-trained models already know via normalisation operations, so that existing tools are more successful on such different data. The other option is to create *native* models by training them on labelled Twitter data. The drawback of the first option is that it's not clear what norm to target: "what is standard language?" (Eisenstein, 2013; Plank, 2016), and implementing normalisation procedures requires quite a lot of manual intervention and subjective decisions. The drawback of the second option is that manually annotated Twitter data isn't readily available, and it is costly to produce.

In this paper, we report on our participation in PoSTWITA[1], the EVALITA 2016 shared task on Italian Part-of-Speech (POS) tagging for Twitter (Tamburini et al., 2016). We emphasise an approach geared to building a *single model* (rather than an ensemble) based on weakly supervised learning, thus favouring (over normalisation) the aforementioned second option of learning *invariant representations*, also for theoretical reasons. We address the bottleneck of acquiring manually annotated data by suggesting and showing that a semi-supervised approach that mainly focuses on tweaking data selection within a bootstrapping setting can be successfully pursued for this task. Contextually, we show that large amounts of manually annotated data might not be helpful if data isn't "of the right kind".

2 Data selection and bootstrapping

In adapting a POS tagger to Twitter, we mainly focus on ways of selectively enriching the training set with additional data. Rather than simply adding large amounts of existing annotated data, we investigate ways of selecting smaller amounts of more appropriate training instances, possibly even tagged with silver rather than gold labels. As

[1] http://corpora.ficlit.unibo.it/PoSTWITA/

for the model itself, we simply take an off-the-shelf tagger, namely a bi-directional Long Short-Term Memory (bi-LSTM) model (Plank et al., 2016), which we use with default parameters (see Section 3.2) apart from initializing it with Twitter-trained embeddings (Section 3.1).

Our first model is trained on the PoSTWITA training set plus additional gold data selected according to two criteria (see below: *Two shades of gold*). This model is used to tag a collection of Facebook posts in a bootstrapping setting with two cycles (see below: *Bootstrapping via Facebook*). The rationale behind using Facebook as *not-so-distant* source when targeting Twitter is the following: many Facebook posts of public, non-personal pages resemble tweets in style, because of brevity and the use of hashtags. However, differently from random tweets, they are usually correctly formed grammatically and spelling-wise, and often provide more context, which allows for more accurate tagging.

Two shades of gold We used the Italian portion of the latest release (v1.3) of the Universal Dependency (UD) dataset (Nivre et al., 2016), from which we extracted two subsets, according to two different criteria. First, we selected data on the basis of its *origin*, trying to match the Twitter training data as close as possible. For this reason, we used the Facebook subportion (`UD_FB`). These are 45 sentences that presumably stem from the Italian Facebook help pages and contain questions and short answers.[2] Second, by looking at the confusion matrix of one of the initial models, we saw that the model's performance was especially poor for cliticised verbs and interjections, tags that are also infrequent in the training set (Table 2). Therefore, from the Italian UD portion we selected *any* data (in terms of origin/genre) which contained the `VERB_CLIT` or `INTJ` tag, with the aim to boost the identification of these categories. We refer to this set of 933 sentences as `UD_verb_clit+intj`.

Bootstrapping via Facebook We augmented our training set with silver-labelled data. With our best model trained on the original task data plus `UD_verb_clit+intj` and `UD_FB`, we tagged a collection of Facebook posts, added those to the training pool, and retrained our tagger. We used two iterations of indelible self-training (Abney, 2007), i.e., adding automatically tagged data where labels do not change once added. Using the Facebook API through the Facebook-sdk python library[3], we scraped an average of 100 posts for each of the following pages, selected on the basis of our intuition and on reasonable site popularity:

- sport: `corrieredellosport`
- news: `Ansa.it`, `ilsole24ore`, `lastampa.it`
- politics: `matteorenziufficiale`
- entertainment: `novella2000`, `alFemminile`
- travel: `viaggiart`

We included a second cycle of bootstrapping, scraping a few more Facebook pages (`soloGossip.it`, `paesionline`, `espressonline`, `LaGazzettaDelloSport`, again with an average of 100 posts each), and tagging the posts with the model that had been re-trained on the original training set plus the first round of Facebook data with silver labels (we refer to the whole of the automatically-labelled Facebook data as `FB_silver`). `FB_silver` was added to the training pool to train the final model. Statistics on the obtained data are given in Table 1.[4]

Table 1: Statistics on the additional datasets.

Data	Type	Sents	Tokens
`UD_FB`	gold	45	580
`UD_verb_clit+intj`	gold	933	26k
`FB` (all, iter 1)	silver	2243	37k
`FB` (all, iter 2)	silver	3071	47k
Total added data	gold+silver	4049	74k

3 Experiments and Results

In this section we describe how we developed the two models of the final submission, including all preprocessing decisions. We highlight the importance of choosing an adequate development set to identify promising directions.

3.1 Experimental Setup

PoSTWITA data In the context of PoSTWITA, training data was provided to all participants in the

[2] These are labelled as `4-FB` in the comment section of UD. Examples include: *Prima di effettuare la registrazione. È vero che Facebook sarà a pagamento?*

[3] https://pypi.python.org/pypi/facebook-sdk

[4] Due to time constraints we did not add further iterations; we cannot judge if we already reached a performance plateau.

Table 2: Tag distribution in the original trainset.

Tag	Explanation	#Tokens	Example
NOUN	noun	16378	cittadini
PUNCT	punctuation	14513	?
VERB	verb	12380	apprezzo
PROPN	proper noun	11092	Ancona
DET	determiner	8955	il
ADP	preposition	8145	per
ADV	adverb	6041	sempre
PRON	pronoun	5656	quello
ADJ	adjective	5494	mondiale
HASHTAG	hashtag	5395	#manovra
ADP_A	articulated prep	4465	nella
CONJ	coordinating conj	2876	ma
MENTION	mention	2592	@InArteMorgan
AUX	auxiliary verb	2273	potrebbe
URL	url	2141	http://t.co/La3opKcp
SCONJ	subordinating conj	1521	quando
INTJ	interjection	1404	fanculo
NUM	number	1357	23%
X	anything else	776	s...
EMO	emoticon	637	😔
VERB_CLIT	verb+clitic	539	vergognarsi
SYM	symbol	334	→
PART	particle	3	's

form of manually labelled tweets. The tags comply with the UD tagset, with a couple of modifications due to the specific genre (emoticons are labelled with a dedicated tag, for example), and subjective choices in the treatment of some morphological traits typical of Italian. Specifically, clitics and articulated prepositions are treated as one single form (see below: *UD fused forms*). The training set contains 6438 tweets, for a total of ca. 115K tokens. The distribution of tags together with examples is given in Table 2. The test set comprises 301 tweets (ca. 4800 tokens).

UD fused forms In the UD scheme for Italian, articulated prepositions (ADP_A) and cliticised verbs (VERB_CLIT) are annotated as separate word forms, while in PoSTWITA the original word form (e.g., 'alla' or 'arricchirsi') is annotated as a whole. In order to get the PoSTWITA ADP_A and VERB_CLIT tags for these fused word forms from UD, we adjust the UCPH ud-conversion-tools[5] (Agić et al., 2016) that propagates head POS information up to the original form.

Pre-processing of unlabelled data For the Facebook data, we use a simplistic off-the-shelf rule-based tokeniser that segments sentences by punctuation and tokens by whitespace.[6] We normalise URLs to a single token (http://www.someurl.org) and add a rule for smileys. Finally, we remove sentences from the Facebook data were more than 90% of the tokens are in all caps. Unlabelled data used for embeddings is preprocessed only with normalisation of usernames and URLs.

Word Embeddings We induced word embeddings from 5 million Italian tweets (TWITA) from Twita (Basile and Nissim, 2013). Vectors were created using word2vec (Mikolov and Dean, 2013) with default parameters, except for the fact that we set the dimensions to 64, to match the vector size of the multilingual (POLY) embeddings (Al-Rfou et al., 2013) used by Plank et al. (2016). We dealt with unknown words by adding a "UNK" token computing the mean vector of three infrequent words ("vip!","cuora", "White").

Figure 1: Word cloud from the training data.

Creation of a *realistic* internal development set The original task data is distributed as a single training file. In initial experiments we saw that performance varied considerably for different random subsets. This was due to a large bias towards tweets about 'Monti' and 'Grillo', see Figure 1, but also because of duplicate tweets. We opted to create *the most difficult* development set possible. This development set was achieved by removing duplicates, and randomly selecting a subset of tweets that do not mention 'Grillo' or 'Monti' while maximizing out-of-vocabulary (OOV) rate with respect to the training data. Hence, our internal development set consisted of 700 tweets with an OOV approaching 50%. This represents a more realistic testing scenario. Indeed, the baseline (the basic bi-LSTM model), dropped from 94.37 to 92.41 computed on the earlier development set were we had randomly selected 1/5 of the data, with an OOV of 45% (see Table 4).

3.2 Model

The bidirectional Long Short-Term Memory model bilty[7] is illustrated in Figure 2. It is a

[5]https://github.com/coastalcph/ud-conversion-tools

[6]https://github.com/bplank/multilingualtokenizer

[7]https://github.com/bplank/bilstm-aux

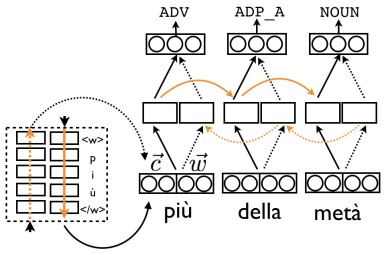

Figure 2: Hierarchical bi-LSTM model using word $\vec{w}$ and character $\vec{c}$ representations.

context bi-LSTM taking as input word embeddings $\vec{w}$. Character embeddings $\vec{c}$ are incorporated via a hierarchical bi-LSTM using a sequence bi-LSTM at the lower level (Ballesteros et al., 2015; Plank et al., 2016). The character representation is concatenated with the (learned) word embeddings $\vec{w}$ to form the input to the context bi-LSTM at the upper layers. We took default parameters, i.e., character embeddings set to 100, word embeddings set to 64, 20 iterations of training using Stochastic Gradient Descent, a single bi-LSTM layer and regularization using Gaussian noise with $\sigma = 0.2$ (`cdim 100, trainer sgd, indim 64, iters 20, h_layer 1, sigma 0.2`). The model has been shown to achieve state-of-the-art performance on a range of languages, where the incorporation of character information was particularly effective (Plank et al., 2016). With these features and settings we train two models on different training sets.

GOLDPICK `bilty` with pre-initialised TWITA embeddings, trained on the PoSTWITA training set plus selected gold data (`UD_FB` + `UD_verb_clit+intj`).

SILVERBOOT a bootstrapped version of GOLDPICK, where `FB_silver` (see Section 2) is also added to the training pool, which thus includes both gold and silver data.

3.3 Results on test data

Participants were allowed to submit one official, and one additional (unofficial) run. Because on development data SILVERBOOT performed better than GOLDPICK, we selected the former for our official submission and the latter for the unofficial one, making it thus also possible to assess the specific contribution of bootstrapping to performance.

Table 3: Results on the official test set. BEST is the highest performing system at PoSTWITA.

System	Accuracy
BEST	93.19
SILVERBOOT (official)	92.25
GOLDPICK (unofficial)	91.85
TNT (on POSTWITA train)	84.83
TNT (on SILVERBOOT data)	85.52

Table 3 shows the results on the official test data for both our models and TNT (Brants, 2000). The results show that adding bootstrapped silver data outperforms the model trained on gold data alone. The additional training data included in SILVERBOOT reduced the OOV rate for the test-set to 41.2% (compared to 46.9% with respect to the original PoSTWITA training set). Note that, on the original randomly selected development set the results were less indicative of the contribution of the silver data (see Table 4), showing the importance of a carefully selected development set.

4 What didn't work

In addition to what we found to boost the tagger's performance, we also observed what didn't yield any improvements, and in some case even lowered global accuracy. What we experimented with was triggered by intuition and previous work, as well as what we had already found to be successful, such as selecting additional data to make up for under-represented tags in the training set. However, everything we report in this section turned out to be either pointless or detrimental.

More data We added to the training data *all* (train, development, and test) sections from the Italian part of UD1.3. While training on selected gold data (978 sentences) yielded 95.06% accuracy, adding all of the UD-data (12k sentences of newswire, legal and wiki texts) yielded a disappointing 94.88% in initial experiments (see Table 4), also considerably slowing down training.

Next, we tried to add more Twitter data from XLIME, a publicly available corpus with multiple layers of manually assigned labels, including POS tags, for a total of ca. 8600 tweets and 160K tokens (Rei et al., 2016). The data isn't provided as a single gold standard file but in the form of

Table 4: Results on internal development set.

System	Accuracy
Internal dev (prior) OOV: 45%	
BASELINE (w/o emb)	94.37
+POLY emb	94.15
+TWITA emb	**94.69**
BASELINE+TWITA emb	
+Morphit! coarse MTL	94.61
+Morphit! fine MTL	94.68
+UD all	94.88
+gold-picked	95.06
+gold-picked+silver (1st round)	**95.08**
Internal dev (realistic) OOV: 50%	
BASELINE (incl. TWITA emb)	92.41
+gold (GOLDPICK)	93.19
+gold+silver (SILVERBOOT)	**93.42**
adding more gold (Twitter) data:	
+XLIME ADJUDICATED (48)	92.58
+XLIME SINGLE ANNOT.	91.67
+XLIME ALL (8k)	92.04

separate annotations produced by different judges, so that we used MACE (Hovy et al., 2013) to adjudicate divergences. Additionally, the tagset is slightly different from the UD set, so that we had to implement a mapping. The results in Table 4 show that adding all of the XLIME data declines performance, despite careful preprocessing to map the tags and resolve annotation divergences.

More tag-specific data From the matrix computed on the dev set, it emerged that the most confused categories were `NOUN` and `PROPN`. Following the same principle that led us to add `UD_verb_clit+intj`, we tried to reduce such confusion by providing additional training data containing proper nouns. This did not yield any improvements, neither in terms of global accuracy, nor in terms of precision and recall of the two tags.

Multi-task learning Multi-task learning (MTL) (Caruana, 1997), namely a learning setting where more than one task is learnt at the same time, has been shown to improve performance for several NLP tasks (Collobert et al., 2011; Bordes et al., 2012; Liu et al., 2015). Often, what is learnt is one main task and, additionally, a number of auxiliary tasks, where the latter should help the model converge better and overfit less on the former. In this context, the additional signal we use to support the learning of each token's POS tag is the token's degree of ambiguity. Using the information stored in *Morph-it!*, a lexicon of Italian inflected forms with their lemma and morphological features (Zanchetta and Baroni, 2005), we obtained the *number* of all different tags potentially associated to each token. Because the *Morph-it!* labels are highly fine-grained we derived two different ambiguity scores, one on the original and one on coarser tags. In neither case the additional signal contributed to the tagger's performance, but we have not explored this direction fully and leave it for future investigations.

5 Conclusions

The main conclusion we draw from the experiments in this paper is that *data selection matters*, not only for training but also while developing for taking informed decisions. Indeed, only after creating a carefully designed internal development set we obtained stronger evidence of the contribution of silver data which is also reflected in the official results. We also observe that choosing less but more targeted data is more effective. For instance, TWITA embeddings contribute more than generic POLY embeddings which were trained on substantially larger amounts of Wikipedia data. Also, just blindly adding training data does not help. We have seen that using the whole of the UD corpus is not beneficial to performance when compared to a small amount of selected gold data, both in terms of origin and labels covered. Finally, and most importantly, we have found that adding little amounts of not-so-distant silver data obtained via bootstrapping resulted in our best model.

We believe the low performance observed when adding xLIME data is likely due to the non-correspondence of tags in the two datasets, which required a heuristic-based mapping. While this is only a speculation that requires further investigation, it seems to indicate that exploring semi-supervised strategies is preferrable to producing idiosyncratic or project-specific gold annotations.

Acknowledgments We thank the CIT of the University of Groningen for providing access to the Peregrine HPC cluster. Barbara Plank acknowledges NVIDIA corporation for support.

References

[Abney2007] Steven Abney. 2007. *Semisupervised learning for computational linguistics*. CRC Press.

[Agić et al.2016] Željko Agić, Anders Johannsen, Barbara Plank, Héctor Martínez Alonso, Natalie Schluter, and Anders Søgaard. 2016. Multilingual projection for parsing truly low-resource languages. *Transactions of the Association for Computational Linguistics (TACL)*, 4:301–312.

[Al-Rfou et al.2013] Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Polyglot: Distributed word representations for multilingual NLP. *arXiv preprint arXiv:1307.1662*.

[Ballesteros et al.2015] Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2015. Improved transition-based parsing by modeling characters instead of words with lstms. In *EMNLP*.

[Basile and Nissim2013] Valerio Basile and Malvina Nissim. 2013. Sentiment analysis on italian tweets. In *Proceedings of the 4th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 100–107.

[Bordes et al.2012] Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2012. Joint learning of words and meaning representations for open-text semantic parsing. In *AISTATS*, volume 351, pages 423–424.

[Brants2000] Thorsten Brants. 2000. Tnt: a statistical part-of-speech tagger. In *ANLP*.

[Caruana1997] Rich Caruana. 1997. Multitask learning. *Machine Learning*, 28:41–75.

[Collobert et al.2011] Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.

[Eisenstein2013] Jacob Eisenstein. 2013. What to do about bad language on the internet. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 359–369, Atlanta.

[Foster et al.2011] Jennifer Foster, Ozlem Cetinoglu, Joachim Wagner, Josef Le Roux, Joakim Nivre, Deirde Hogan, and Josef van Genabith. 2011. From news to comments: Resources and benchmarks for parsing the language of Web 2.0. In *IJCNLP*.

[Gimpel et al.2011] Kevin Gimpel, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. Part-of-Speech Tagging for Twitter: Annotation, Features, and Experiments. In *Proceedings of ACL*.

[Hovy et al.2013] Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. Learning whom to trust with MACE. In *NAACL*.

[Liu et al.2015] Xiaodong Liu, Jianfeng Gao, Xiaodong He, Li Deng, Kevin Duh, and Ye-Yi Wang. 2015. Representation learning using multi-task deep neural networks for semantic classification and information retrieval. In *Proc. NAACL*.

[McClosky et al.2010] David McClosky, Eugene Charniak, and Mark Johnson. 2010. Automatic domain adaptation for parsing. In *NAACL-HLT*.

[Mikolov and Dean2013] T Mikolov and J Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*.

[Nivre et al.2016] Joakim Nivre et al. 2016. Universal dependencies 1.3. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Charles University in Prague.

[Plank et al.2016] Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss. In *ACL*.

[Plank2016] Barbara Plank. 2016. What to do about non-standard (or non-canonical) language in NLP. In *KONVENS*.

[Rei et al.2016] Luis Rei, Dunja Mladenic, and Simon Krek. 2016. A multilingual social media linguistic corpus. In *Conference of CMC and Social Media Corpora for the Humanities*.

[Tamburini et al.2016] Fabio Tamburini, Cristina Bosco, Alessandro Mazzei, and Andrea Bolioli. 2016. Overview of the EVALITA 2016 Part Of Speech on TWitter for ITAlian Task. In Pierpaolo Basile, Franco Cutugno, Malvina Nissim, Viviana Patti, and Rachele Sprugnoli, editors, *Proceedings of the 5th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2016)*. aAcademia University Press.

[Zanchetta and Baroni2005] Eros Zanchetta and Marco Baroni. 2005. Morph-it! A free corpus-based morphological resource for the Italian language. *Corpus Linguistics 2005*, 1(1).